

Beam simulations for IRE and driver — status and strategy¹

Alex Friedman^{a,b,*}, David P. Grote^{a,b}, Edward P. Lee^{a,c}, Eric Sonnendruker^{c,d}

^a*Heavy Ion Fusion Virtual National Laboratory*

^b*Lawrence Livermore National Laboratory, University of California*

^c*E. O. Lawrence Berkeley National Laboratory, University of California*

^d*CNRS, IRMA Strasbourg and IECN Nancy*

Abstract

The methods and codes employed in the U.S. Heavy Ion Fusion program to simulate the beams in an Integrated Research Experiments (IRE) facility and a fusion driver are presented in overview. A new family of models incorporating accelerating module impedance, multi-beam, and self-magnetic effects is described, and initial WARP3d particle simulations of beams using these models are presented. Finally, plans for streamlining the machine-design simulation sequence, and for simulating beam dynamics from the source to the target in a consistent and comprehensive manner, are described.

PACS classification codes:

29.27.-a Beams in particle accelerators

52.65.Cc Particle orbit and trajectory

52.65.Rr Particle-in-cell method

Keywords:

Accelerator, Fusion, Heavy-ion, Induction, Simulation, Particle-in-cell, Plasma, Beam

*Corresponding author and address:

Dr. Alex Friedman

Mail Stop 47-112, Lawrence Berkeley National Laboratory

One Cyclotron Road, Berkeley, CA 94720

email: af@llnl.gov tel: 510-486-5592 fax 510-486-5392

¹This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore and Lawrence Berkeley National Laboratories under Contracts No. W-7405-ENG-48 and DE-AC03-76SF00098.

I. Introduction

Computer simulations have played an important role in the U.S. Heavy Ion Fusion (HIF) program from its outset. Improved methods and codes are being employed to simulate the beams in an Integrated Research Experiments (IRE) facility and a fusion driver, and are presented in overview, in Section II, below. A new family of models incorporating accelerating module impedance, multi-beam, and self-magnetic effects, and initial WARP3d particle simulations of beams using these models, are presented in Section III. Finally, plans for streamlining the machine-design simulation sequence, and for simulating beam dynamics from the source to the target in a consistent and comprehensive manner, are described in Section IV.

This paper is by no means a complete summary of recent work. A first study assessing the requirements for beam steering in the presence of machine errors in a “model” IRE, conducted using WARPxy, and initial 3-D simulations of the electrostatic-focusing section of that design, conducted using WARP3d, were presented as part of the first author’s talk at the Symposium. This work is described in [1,2]. Work is also in progress to assess the required tolerances in accelerating and “ear” waveforms, magnet fields, alignment, and other aspects of the machine, and to evaluate final focusing optics designs for IRE and driver.

II. Overview of Methods and Codes

Beam simulations fall into three general classes: those that follow a representative set of particles via a Monte-Carlo method known as particle-in-cell (PIC); those that evolve the beam distribution function $f(\mathbf{x}, \mathbf{v})$ in time (nonlinear-perturbative f and semi-Lagrangian Vlasov methods); and those that evolve moments of f . In addition, a zero-dimensional systems model based on scaling laws is used for overall design. Most simulations for Heavy Ion Fusion (HIF) have been carried out using PIC methods, but the other methods have merits that make them the tools of choice for some applications.

Particle-in-cell codes

The principal PIC code used for accelerator studies is WARP, which combines features of a plasma simulation code and an accelerator code. Several geometries are available. WARP3d employs a “warped Cartesian” 3D computational mesh (to accommodate bends in the beam line) and a 6D phase space wherein each particle is characterized by its (x, y, z, p_x, p_y, p_z) . WARPxy uses a transverse-slice (steady flow) description with a 2D mesh and a 5D phase space wherein each particle is characterized

by its (x, y, p_x, p_y, p_z) ; here the independent coordinate is effectively path length along the (possibly bent) system axis. WARPrz uses an axisymmetric description with a 2D (r, z) mesh and a 5D phase space $(r, z, p_r, p_z, p_\theta)$. In all geometries the mesh on which the self-consistent field is computed is a moving window (it may remain fixed in the laboratory frame until the beam has been fully injected, then begin tracking the beam); at present the description is electrostatic, but a Darwin (magnetoinductive) model is planned (see the final section). Descriptions of the “lattice” of focusing, bending and accelerating elements are available at varying levels of detail, including a sharp-edged model, a one-dimensional table of multipole moments vs. axial position, a 3D grid, and first-principles treatment of internal electrostatic elements (including subgrid-scale placement of their boundaries to avoid staircasing). WARP offers a user-programmable interactive interpreter interface based on the freely-available Python language, and employs both parallelization and improved algorithms to enhance run speed. The code is further described in [3]; see also [1].

For studies of beams in the fusion chamber, three PIC codes are in use. All employ a fully electromagnetic field description; a Darwin model would not, in general, afford larger timesteps, since the electron plasma frequency and not the light wave transit time across a cell sets the step size. The first of these codes, BICrz, employs a spatially converging mesh to preserve resolution of the beam as its radius diminishes [4]. This code has been used to study neutralized-ballistic propagation. A newer code, BPIC, offers both r, z and 3D geometries [5]. Its transverse mesh size is variable in time, rather than position, with a special algorithm to advect away the attendant errors. An advanced version of this code offers overlaid fine and coarse meshes, so that, with further development, first-principles treatments of inter-beam effects may be studied. The third code, LSP, also offers both r, z and 3D geometries [6]. In addition, it employs an implicit advance, and a hybrid fluid/particle electron model, so that much denser plasmas may be studied without incurring severe timestep-size limitations. Thus, it is ideal for treating various high-density chamber transport modes. However, LSP does not yet offer a converging mesh, and so carefully tapered transverse zone sizes must be used. These tools are currently being benchmarked against each other.

Codes which evolve a distribution function

The nonlinear-perturbative, or f , method evolves the perturbation f to the distribution function f along particle orbits. When f is small (in the sense of an integral over velocities) compared with the equilibrium f_0 , the statistical noise is effectively

reduced by the ratio of the two, relative to conventional PIC. This makes the method a good choice for examining detailed mode behavior, since eigenfunctions and growth rates are cleanly observable, and for studying the slow evolution of a beam in a storage ring. This class of methods is less ideal for long-time simulation of a beam in an induction linac (since the beam evolves far from its initial state), and for other situations in which the perturbation to the distribution function becomes large or would be represented coarsely by very few particles. The f method is embodied in BEST, which is being applied to studies of beams in both the driver and the fusion chamber [7]; BEST can also run in a PIC mode.

It is also becoming feasible to evolve the values of f on the nodes of a Cartesian mesh in the 4D phase space (x, y, p_x, p_y) ; full 6D will eventually become practical. While several algorithms exist, a promising choice is the semi-Lagrangian Vlasov method, as embodied in the SLV code [8,9]. In this method, the calculation reaches backward in time along a characteristic (orbit in phase space) to obtain the current value of f at each node. Thus, low-density regions of phase space are tracked with the same accuracy as high-density regions; this is useful for halo studies. This method naturally coarse-grains the phase space on the scale of the computational mesh. The errors it introduces are diffusive, and differ in character from those of PIC, making comparisons useful.

Codes which follow moments of the distribution

The principal code in this category is CIRCE, which uses a Lagrangian cold fluid model (with hundreds of discrete slices) to describe the longitudinal dynamics, along with moment equations for the transverse centroid and envelope extent of each slice [10]. Due to the code's speed, it is used heavily for a number of applications, including synthesis of acceleration and compression schedules, transport lattice improvement, assessment of tolerances for accelerating and “ears” waveforms, studies of alignment tolerances, beam sensing, and steering, and studies of drift compression and pulse shaping. The CIRCE model has the following limitations: there is no model for emittance growth or phase-mixing of “mismatch,” slow variation of quantities along beam is assumed, and there are no module impedance, self-magnetic, or inductive models. Some of these limitations are not fundamental.

III. Module Impedance Effects on Longitudinal Dynamics

A schematic of an accelerating module is shown in Fig. 1. Note that the high energy part of a driver consists mostly of such modules (“gaps”) ; we denote the gap length by l_g

and the insulator radius by r_g . We have implemented into WARP3d a simple model for the impedance of the module as it affects the beam. The pulse-forming line absorbs energy from the beam as a beam-induced wave propagates up the line; the line is represented as a resistance R in parallel with the capacitance C of the induction module. The circuit is depicted in Fig. 2. To date we have neglected the core impedance (which is mostly resistive, and sub-dominant) in our simulations. That impedance is associated with the material and format of the ferromagnetic cores (tightly-wound spirals of amorphous metallic glass tape with inter-laminar insulation); a complete model has yet to be developed. The circuit equation is:

$$RC \frac{dV}{dt} + V = V_{\text{ext}} - I_b R \quad (1)$$

For a drifting beam, this becomes:

$$RC \frac{dV}{dt} + V = - (I_b - I_{b0}) R \quad (2)$$

where I_{b0} is the time-shifted initial beam current. The beam ends must be confined against their own thermally- and space charge-induced axial expansion and, to this end, confining “ears” are included in the imposed voltage waveforms, along with the main accelerating and pulse-compressing components. These ear fields are assumed to be generated by separate pulsers with small R and C . We compute I_b as an average of $I_b(z)$ over the length of the gap, and find that a simple backward difference suffices (a centered difference causes difficulty when $C = 0$).

To begin to explore the effects of module impedance on longitudinal dynamics, we began with simulations of a 10 GeV, 3 kA drifting beam in the limit of $R = C = 0$; these parameters correspond to a scenario (such as an induction recirculator) employing a small number of beams, and were chosen to make contact with earlier work [11]. We applied a perturbing “bump” to the velocity distribution at mid-pulse. The results of such a simulation are shown in Fig. 3, which depicts the line charge density versus time at a set of “stations” down the beam line; the time $t = 0$ at each station coincides with the arrival of the leading edge of the computational grid at that station. Successive curves are offset in ordinate from each other so as to render the wave characteristics visible. Wave reflection at the beam ends is clearly visible. In this run, and in those shown in Fig. 4, the beam travels through 3 km, 500 lattice periods, and 1000 gaps in $24.6\mu\text{s}$, using 75,000 time steps. Also, $A_{\text{ion}} = 130$, $v_b = 1.2 \times 10^8$ m/s, $l_b = 10.8$ m, phase advance per lattice period $\phi_0 = 70^\circ$ depressed by space charge to $\phi = 15^\circ$, beam semi-axes $a_0 = 3.2$ cm and $b_0 = 1.8$ cm, a square metal pipe at $x_w = 5$ cm, half-lattice period length $l_{\text{hlp}} = 3$ m, 150 steps per period, and grid sizes $z = 2.34$ cm, $x = y = 6.25$ mm.

The runs shown here used a very crude (8x8) transverse mesh (for a quadrant), 512 cells in z, and 40,000 particles, and required a few hours on a workstation. Full 3-D runs using 640,000 particles and a 32x32x512 mesh were carried out as a check; those required 1.77 hours on 128 processors of a Cray-T3E computer. The results (for the quantities shown here) were very similar, but differed in detail. When it is important to quantitatively capture both transverse and longitudinal physics, such runs are required.

When a resistance of 600 Ohms is incorporated in a simulation wherein a random velocity perturbation (a function of axial position) was applied, waves which propagate in a backward direction on the beam are seen to be unstable. Figure 4(a) shows the perturbation to the line charge density as it evolves from station to station; here the nominal line charge density has been subtracted out for clarity; the intersection of each curve with the ordinate axis denotes that observing stations' location, rather than (as in the previous figure) the magnitude of the perturbation; in this figure the perturbed line charge density has been arbitrarily scaled for visibility. In Fig. 4(b), a module capacitance $C = 0.033$ nF has been included in the simulation, and its stabilizing effect is evident. Analytic theory does not predict complete stabilization for a cold beam; we conjecture that the finite longitudinal thermal spread also contributes to the observed stability.

A driver is expected to employ a larger number of beams each with smaller current, and simulations of that regime began with single-beam studies. Such a run is shown in Fig. 5, which is for a Cs^+ beam at 1.76 GeV perturbed at the start of the run by a velocity “bump” at mid-pulse. Some parameters for these runs are: $I_b = 50$ A, $l_{\text{hlp}} = 3.91$ m, $l_g = 3$ m, $a_0 = 1.44$ cm, initial beam semi-axes 1.97 cm and 1.05 cm, $x_w = 3.44$ cm (square pipe), $A_{\text{ion}} = 133$, $v_b = 5 \times 10^7$ m/s, $\beta = .167$, $l_b = 15$ m, $R_{\text{gap}} = 782$ Ω , $C_{\text{gap}} = 0$, $\theta_0 = 75^\circ$, $\theta = 20^\circ$, $V_{\text{thermal,z}} = .5$ km/s, $\omega_p = 1.12 \times 10^7$ s⁻¹, $\omega_0 = 8.37 \times 10^6$ s⁻¹, $n = 9.55 \times 10^{15}$ m⁻³, $t = 1.04$ ns, steps/period = 150, $z = 1.53$ cm, $x = 4.3$ mm, 80,000 particles, 8x8x1024 cells/quadrant. With a resistive module impedance, damping of the “forward” wave, and growth of the “backward” wave, is clearly evident.

IV. Multi-Beam and Self-Magnetic Effects

The shielding plates shown in Fig. 1 periodically short out the radial electric field and so limit inter-beam space charge and transverse beam deflections. In cylindrical (r, θ , z) coordinates, and in the limit of plates closely spaced along the axis so that on the large (multi-beam array) scale $E_r = 0$, Faraday's law becomes:

$$\frac{E_z}{r} = \frac{B}{t} \quad (3)$$

The field equation is linear, and so we can superpose the solutions of two systems to solve the problem of interest: the first with the set of N beams each having current I_b inside a metallic outer wall at the insulator radius r_g obeying potential $\phi = 0$, and the second having no beam but accounting for the fields induced in the cavity by the accelerating module in response to the return current. Integrating from an arbitrary radius to r_g (where $E_z = 0$), we obtain:

$$E_{z,\text{ind}}(r) = - \int_r^{r_g} dr' \frac{B(r')}{t} = - \int_r^{r_g} dr' \frac{\mu_0}{2} \frac{N I_b r'^2}{r_g^2} \frac{N v_b}{4 \epsilon_0 c^2} \left(1 - \frac{r'^2}{r_g^2} \right) \frac{I_b}{z} \quad (4)$$

$$- g_{\text{ind}} \frac{b}{z} \quad \text{where} \quad g_{\text{ind}} = - \frac{N^2}{4 \epsilon_0} \left(1 - \frac{r^2}{r_g^2} \right)$$

The electrostatic longitudinal field for an “incompressible” beam is approximated by:

$$E_{z,\text{es}} = - g_{\text{es}} \frac{b}{z} \quad \text{where} \quad g_{\text{es}} = \frac{1}{4 \epsilon_0} 2 \ln \frac{b}{a} \quad (5)$$

where a is the beam radius and b an effective single-beam pipe radius; this form is not used explicitly, since $E_{z,\text{es}}$ is obtained via the usual solution of Poisson’s equation in WARP3d. Thus the inductive field effect is strongest for the central beam(s); on average, for a gap “residency factor” (fraction of the half-lattice period) f_{gap} ,

$$\frac{E_{z,\text{ind}}}{E_{z,\text{es}}} = - \frac{N^2 f_{\text{gap}}}{2 \ln(b/a)} \quad (6)$$

where a is the beam radius and b an effective single-beam pipe radius. Thus, to compute this inductive field in the vicinity of the central beam (where it is strongest), we carry out a single-beam simulation and then multiply the return current by N before using it in the circuit equation, and use the longitudinal derivative of the beam current as computed by WARP (and smoothed over a few grid cells to minimize computational noise):

$$E_{z,\text{ind}}(\text{central beam}) = \frac{N v_b}{4 \epsilon_0 c^2} \frac{I_b}{z} \quad (7)$$

The beams are isolated from each other when they are in the quadrupole magnets, so we apply this field only to those parts of the simulated beam that are in the gap.

Figure 6 shows the contributions to the beam’s E_z field versus position at the tail of the beam, for the central beam in an array of 50 Cs^+ beams at 1.76 GeV. Other parameters are as for the run shown in Fig. 5. Here, $z = 0$ is the left (tail) end of WARP3d’s computational mesh; the jaggedness in the electrostatic field is a result of particle statistics. In a simulation incorporating the inductive E_z model, we find that the ear fields

must compensate for both the electrostatic and inductive beam fields if the beam ends are to be smoothly confined; use of ears adjusted for electrostatic forces alone causes the launching of large waves on the beam. Figure 7 shows the longitudinal phase space after 977.5 m (19.55 μ s) for two simulations (with zero module impedance) of the central beam perturbed at the start of the run by a velocity “bump” at mid-pulse, with and without inductive forces and ear corrections. The vertical fiducials make it clear that the wave speed is larger in the run without inductive E_z by a ratio of 1.5; theory predicts a speed ratio of $(g_{es}/g_{net})^{1/2} = 1.6$, where $g_{net} = g_{es} + g_{ind}$.

For an edge beam, the inductive E_z is negligible, but the magnetic field of the array of beams introduces a transverse “pinching” force (in addition to the electrostatic “bulging”) which is largest for a beam at the edge of the array:

$$B_{,edge} = \frac{\mu_0}{2} \frac{NI_b}{r_{gap}} \quad (8)$$

As for the longitudinal force, when the plates are close enough together and there are enough beams this magnetic force will dominate over the electrostatic bulging. The electrostatic bulging force in general will not vary from beam to beam linearly with the radius of its centroid, but (due to the shielding) will rise sharply as the edge of the array is neared. Thus, one might adjust the plate spacing so that the two forces cancelled for beams at or near the edge of the array, but global cancellation may not be achievable.

V. Plans and Discussion

The process by which a machine design is developed and simulated requires at present multiple steps, some of which involve manual input see Fig. 8(a). One goal is to streamline this process, leading to a sequence similar to the one shown in Fig. 8(b). The “afterburner” referred to is coding which takes as input the overall systems design produced by the IBEAM systems code, and yields a discrete design with, *e.g.*, an integer number of beams and the locations and strengths of the individual quadrupole magnets. At this writing, the integration of CIRCE physics into WARP is partially complete.

Our long-term goal is integrated source-to-target simulation of the beams in an IRE and driver; see Fig. 9. This will require linking particle and field data from the driver simulations, into the chamber propagation simulations, and ultimately as beam-cluster data into the radiation-hydrodynamics calculations used for target design. In addition, linkages to other codes (using a variety of models, as discussed in this paper) will be implemented. These will be used for detailed simulations of sections of the machine; the

linkages will ensure consistent initial conditions for simulations that begin with an evolved beam.

Considerable work is planned in the general area of multi-beam and module impedance effects. By increasing the number of beams, it is possible to reach a regime where $g_{\text{net}} < 0$. We plan to explore that regime in the near future. We also plan to implement the transverse forces described in Section IV into WARP, and to study beam deflections. One can conceive of carrying out a set of CIRCE calculations for all beams, using emittance data from a single beam, to accurately compute deflections (the CIRCE model would need to be extended to handle self-inductive effects, as described below) [12]. In addition, it is possible to extend the field model in WARP3d to account for a variety of physics effects from nearly first principles. One can combine a detailed simulation of a single beam (in electrostatic or Darwin approximation) with field solutions over a larger multi-beam domain (assuming the beams behave similarly), relaxing the approximation of closely-spaced plates made in Section IV. A formalism for carrying this out has been developed [13].

Finally, as noted above, the Darwin model may prove useful in several contexts. That description omits electromagnetic radiation and retardation effects, but is valid for low- to moderate-frequency behavior in an HIF driver. Straightforward time-differencing of the Darwin equations, unstable in most other contexts [14], is expected to be stable in many HIF beam contexts when only ions are to be included, since the stability condition on a grid of overall length L is: $v_p L / c < 1$. Other algorithmic simplifications are also possible for beam applications; see [15] and references therein.

Acknowledgments

The authors gratefully acknowledge contributions to both the presentation and this paper by John Barnard, Bill Herrmannsfeldt, Irving Haber, Rami Kishek, Christine Celata, Michiel de Hoon, George Craig, Steve Lund, Bill Sharp, Wei-Li Lee, and Wayne Meier.

References

- [1] D. P. Grote, *et. al., these Proceedings*.
- [2] J. J. Barnard, *et. al., these Proceedings*.
- [3] D. P. Grote, A. Friedman, I. Haber, W. Fawley, J. L. Vay, “New Developments in WARP3d: Progress Toward End-to-End Simulation,” *Nuclear Inst. and Methods A* **415**, 428 (1998).

- [4] W. M. Sharp, *et. al.*, *these Proceedings*; D. A. Callahan, *Fusion Eng. And Des.* **32-33**, 441 (1996); A. B. Langdon, *Part. Accel.* **37-38**, 175 (1992).
- [5] J-L. Vay and C. Deutsch, *these Proceedings*.
- [6] D. R. Welch, *et. al.*, *these Proceedings*; D. V. Rose, *et. al.*, *these Proceedings*
- [7] H. Qin, *et. al.*, *these Proceedings*
- [8] E. Sonnendrucker, J. R. Roche, P. Bertrand, A. Ghizzo, "The Semi-Lagrangian Method for the Numerical Resolution of the Vlasov Equation," *J. Comput. Phys.* **149**, 201 (1999).
- [9] E. Sonnendrucker, *et. al.*, *these Proceedings*.
- [10] W. M. Sharp, D. P. Grote, and D. L. Judd, "Shaping and Focusing of High-Current Pulses for Heavy-Ion Fusion," paper TUP145, *Proc. 1999 APS/IEEE Particle Accelerator Conference*.
- [11] D. A. Callahan, A. B. Langdon, A. Friedman, and I. Haber, "The Longitudinal Wall Impedance Instability in a Heavy-Ion Fusion Driver," *J. App. Phys.* **81**, 2298 (1997).
- [12] C. M. Celata, LBNL, *private communication* (1999).
- [13] E. Sonnendrucker, A. Friedman, and E. P. Lee, "A model for incorporating module impedance and multi-beam effects into WARP," *LBNL Heavy Ion Fusion note* HIFAN 1001 (1999); A. Friedman, D. A. Callahan-Miller, D. P. Grote, E. P. Lee, and E. Sonnendrucker, "Models for simulating driver-scale beams in accelerators for Heavy Ion Fusion," *presentation at APS Division of Plasma Physics Meeting*, Seattle, Nov. 1999; both available from the authors.
- [14] C. E. Nielson and H. R. Lewis, *Methods in Comput. Phys.* **16**, 367 (1976).
- [15] W. W. Lee *et. al.*, *these Proceedings*.

Figure Captions

Fig. 1. Geometry of a “typical” accelerating module.

Fig. 2. Circuit model employed in simulation studies of module impedance effects.

Fig. 3. 3-D WARP simulation of 10 GeV, 3 kA drifting beam, perturbed at mid-pulse by a velocity "bump." Beam travels through 3 km; 500 lattice periods; 1000 gaps in 24.6 μ s, using 75,000 timesteps. Accelerating module $R = C = 0$.

Fig. 4. Simulations beginning with random perturbation to longitudinal velocity: (a) with module $R = 600$ Ohms and $C = 0$; (b) with $R = 600$ Ohms, $C = 0.033$ nF.

Fig. 5. Evolution of line-charge density for 50 A single-beam simulation (see text).

Fig. 6. Longitudinal fields at beam end for central beam of 50-beam array (see text).

Fig. 7. Longitudinal phase space after 19.5492 μ s (977.5 m): (a) no inductive self-forces or ears correction; (b) with inductive self-forces and ears correction. Vertical lines are fiducials.

Fig. 8. Present-day (a) and envisioned (b) work flow for systems design and simulation.

Fig. 9. Schematic overview of a source-to-target simulation capability.




Fig. 1. Geometry of a “typical” accelerating module

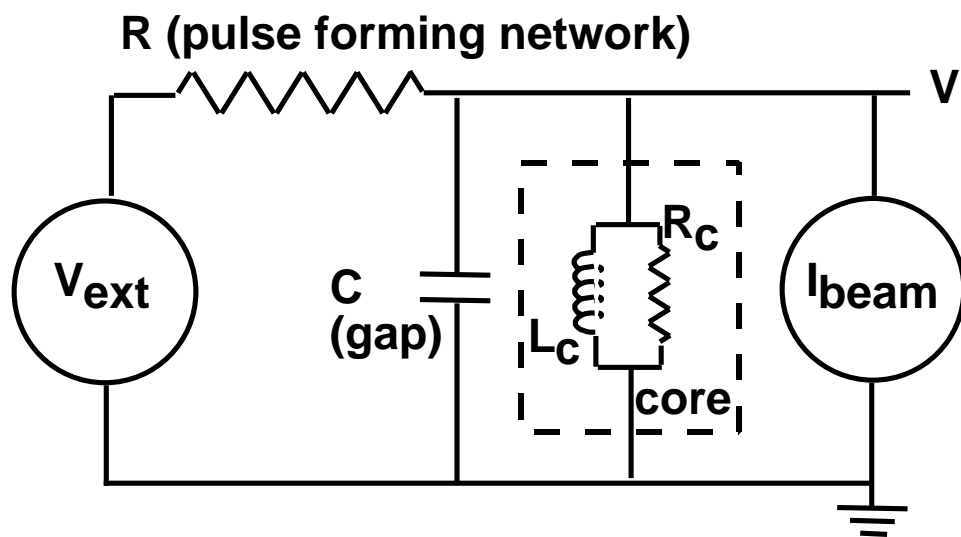


Fig. 2. Circuit model employed in simulation studies of module impedance effects.




Fig. 3. 3-D WARP simulation of 10 GeV, 3 kA drifting beam, perturbed at mid-pulse by a velocity "bump." Beam travels through 3 km; 500 lattice periods; 1000 gaps in $24.6\mu\text{s}$, using 75,000 timesteps. Accelerating module $R = C = 0$

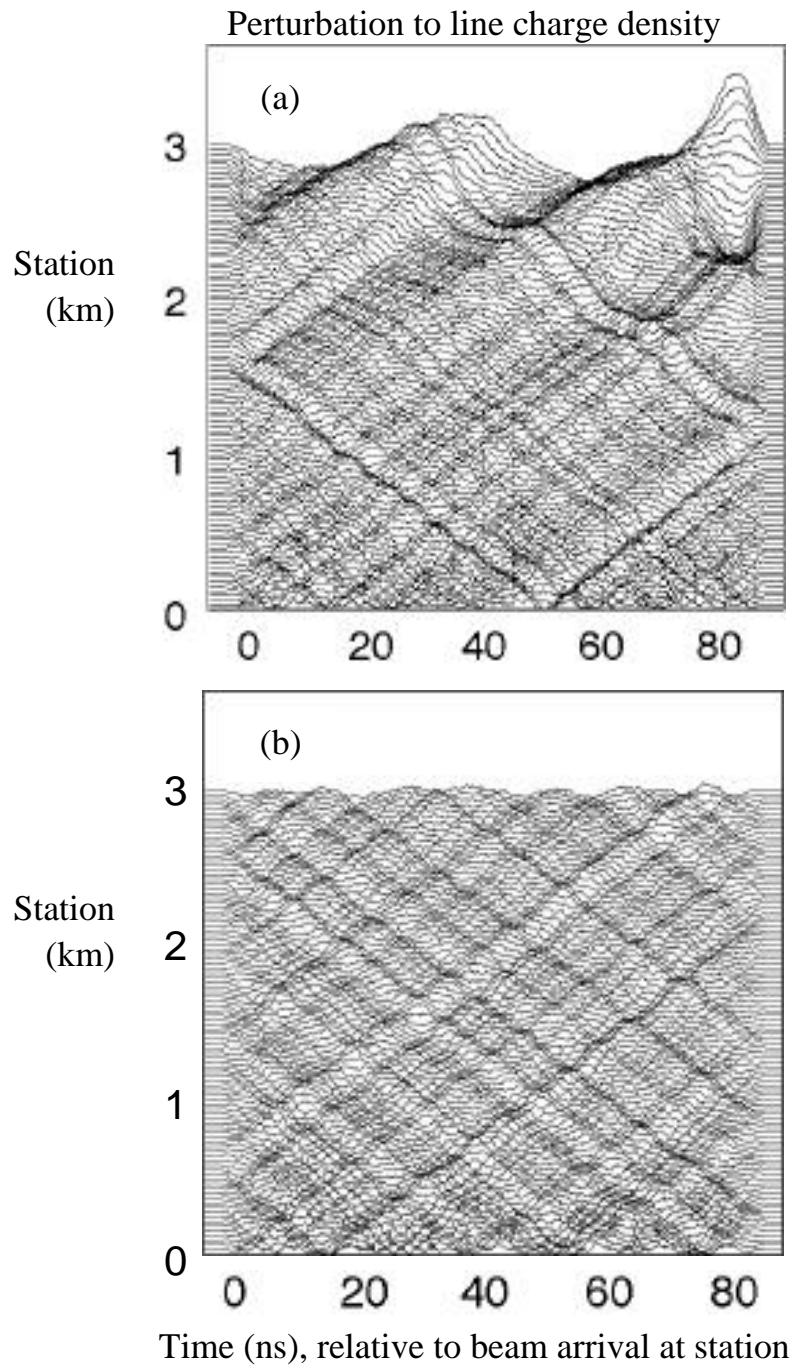


Fig. 4. Simulations beginning with random perturbation to longitudinal velocity: (a) with module $R = 600$ Ohms and $C = 0$; (b) with $R = 600$ Ohms, $C = 0.033$ nF.




Fig. 5. Evolution of line-charge density for 50 A single-beam simulation (see text).




Fig. 6. Longitudinal fields at beam end for central beam of 50-beam array (see text).

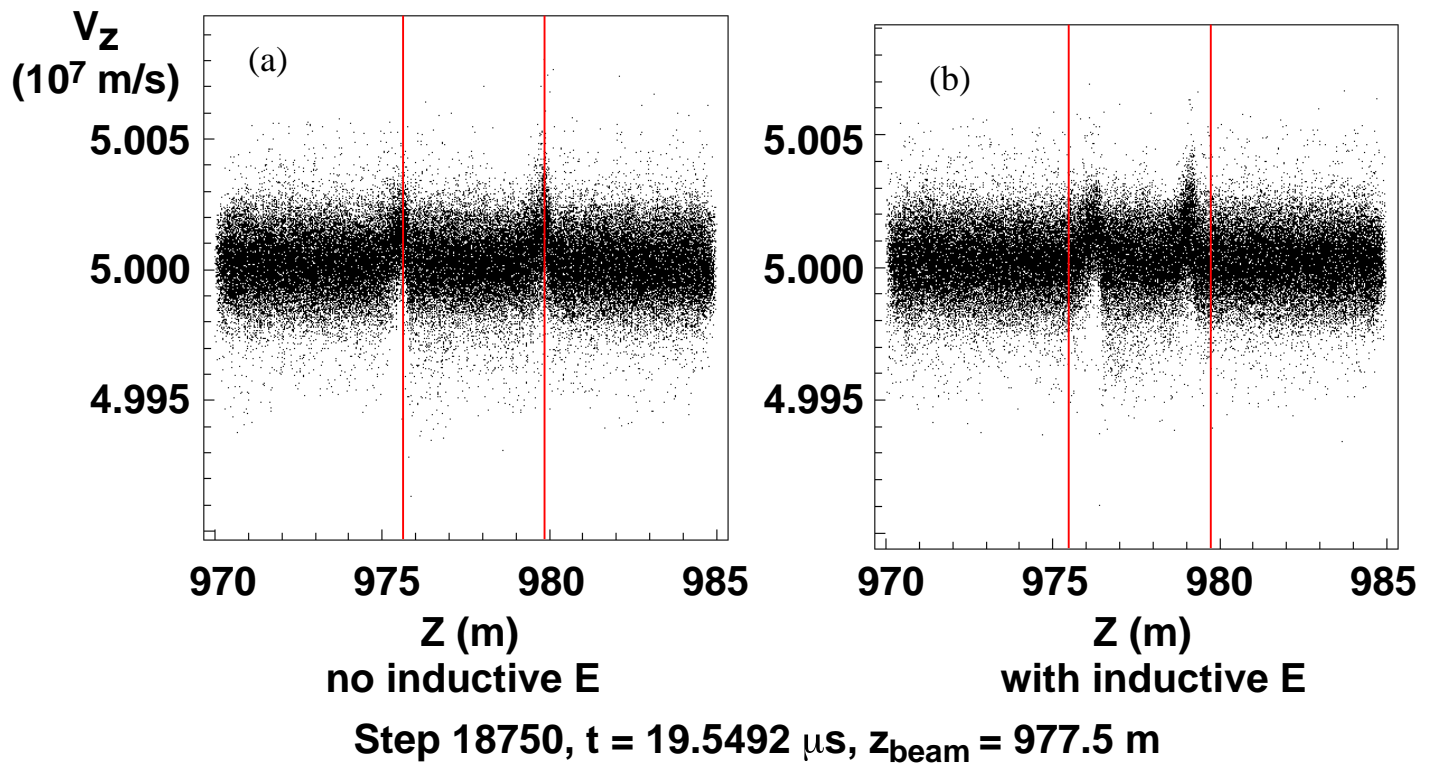


Fig. 7. Longitudinal phase space after $19.5492 \mu\text{s}$ (977.5 m): (a) no inductive self-forces or ears correction; (b) with inductive self-forces and ears correction. Vertical lines are fiducials.

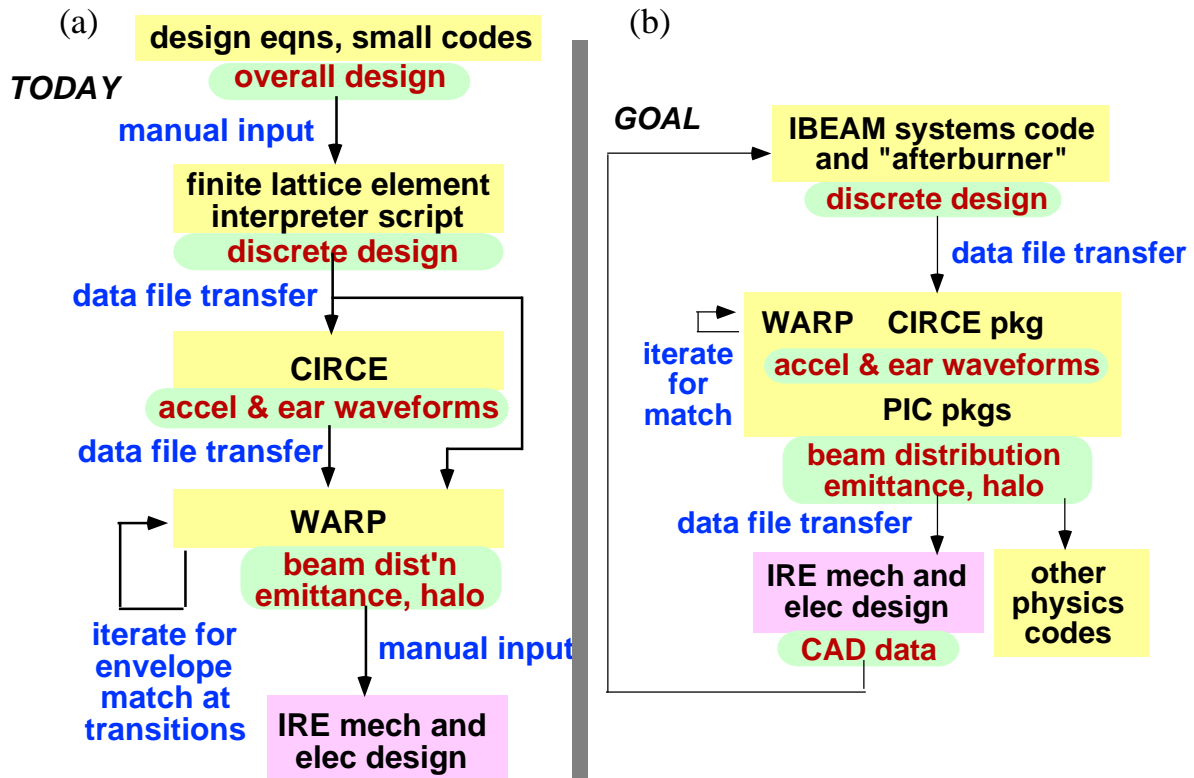


Fig. 8. Present-day (a) and envisioned (b) work flow for systems design and simulation.




Fig. 9. Schematic overview of a source-to-target simulation capability.